

Diagnosing Grey Areas in National Green Taxonomy Frameworks and Corporate Misleading Statements

Authors

Abbas Ahsun, Abill Robert, Abilly Elly

Date; 24 October 2026

Abstract

National green taxonomy frameworks have emerged as the central regulatory instrument for distinguishing genuinely sustainable economic activities from those that merely claim to be so. Yet a persistent diagnostic gap remains: taxonomies classify activities, not corporate intent, and firms routinely exploit the space between technical eligibility and substantive alignment to construct misleading environmental narratives. This study addresses that gap by developing and validating a diagnostic framework for identifying grey areas in taxonomy-based corporate disclosures and linking them to patterns of misleading statements. Drawing on a mixed-methods design combining content analysis of taxonomy-aligned disclosures from 240 European non-financial firms (2022–2024) with a structured expert validation panel, the study constructs a four-dimensional diagnostic instrument encompassing eligibility-alignment divergence, capital expenditure denominator manipulation, narrative-verification asymmetry, and cross-objective double counting. Results demonstrate that the framework achieves 89.4% classification accuracy in identifying firms whose taxonomy disclosures exhibit material divergence from underlying financial and operational data, with capital expenditure denominator errors and selective activity reporting emerging as the strongest predictors of misleading statement risk. The study concludes that taxonomy frameworks require complementary diagnostic instruments to function as credible

anti-greenwashing mechanisms, and offers a replicable tool for regulators, auditors, and investors seeking to distinguish substantive alignment from performative compliance.

Keywords: green taxonomy, greenwashing diagnostics, corporate environmental disclosure, capital expenditure alignment, sustainability governance

1. Introduction

1.1 Background

The proliferation of national and regional green taxonomy frameworks represents one of the most ambitious attempts to impose analytical order on the contested terrain of corporate environmental claims. The European Union Taxonomy, operational since 2022, establishes technical screening criteria across six environmental objectives, requiring firms to report the proportion of turnover, capital expenditure, and operational expenditure that is either “eligible” under the taxonomy or actually “aligned” with its substantive requirements . Similar frameworks have emerged globally, from the UK Green Taxonomy to South Korea’s K-Taxonomy, each promising to transform vague sustainability rhetoric into verifiable metrics .

The regulatory logic is seductively straightforward: if activities can be classified, claims can be verified; if claims can be verified, greenwashing can be curtailed. This logic has driven substantial institutional investment. The Luxembourg financial regulator’s 2025 review of twenty non-financial issuers found that while all used mandated reporting templates, only half provided meaningful explanations for year-on-year changes in taxonomy key performance indicators, and forty percent had errors in capital expenditure denominator calculations . These findings suggest that formal compliance with taxonomy reporting requirements coexists with substantive opacity—a condition this study terms the “grey area” of taxonomy governance.

Parallel scholarship has documented the mechanisms through which corporate environmental communication diverges from operational reality. Okeke (2026) demonstrates that greenwashing operates through four interrelated dimensions—misleading environmental claims, strategic disclosure practices, ESG metric manipulation, and green trust erosion—and that these dimensions can be reliably measured through a validated diagnostic instrument . Okeke and Ugbebor (2026) further theorise corporate sustainability discourse as performative, arguing that purpose claims acquire governance significance not through their accuracy but through their discursive configuration and temporal persistence . These contributions establish that misleading statements are not aberrations but structured features of sustainability communication, requiring equally structured diagnostic responses.

What remains underdeveloped is the integration of these insights with the specific institutional architecture of national taxonomy frameworks. Taxonomies create new categories of claims—

alignment percentages, eligible activity codes, capital expenditure plans—that did not exist in earlier greenwashing contexts. The diagnostic tools developed for evaluating generic ESG disclosures may not capture the distinctive vulnerabilities created by taxonomy reporting: the denominator manipulation opportunities identified by the CSSF, the double-counting risks across multiple environmental objectives, and the performative use of “alignment” language that technical eligibility enables .

1.2 Problem Statement

Existing approaches to detecting corporate environmental misrepresentation fall into three broad categories, each with significant limitations when applied to taxonomy-based disclosures. First, regulatory compliance reviews focus on whether mandated templates are completed and whether numerical entries reconcile with financial statements. The CSSF review exemplifies this approach, identifying denominator errors and missing references but not systematically assessing whether firms are using taxonomy categories strategically to construct misleading impressions . Compliance reviews answer the question “did the firm report correctly?” but not “is the firm’s reporting misleading?”

Second, academic greenwashing measurement instruments, including the validated scale developed by Okeke (2026), capture perceptual and reputational dimensions of misleading communication but were designed for general sustainability claims rather than taxonomy-specific disclosure architectures . The four dimensions of Okeke’s instrument—misleading claims, strategic disclosure, metric manipulation, and green trust—map imperfectly onto the distinctive features of taxonomy reporting, where the manipulation occurs not primarily through false statements but through selective activity classification, denominator construction, and the exploitation of the eligibility-alignment distinction.

Third, firm-level typologies based on ESG score and taxonomy capital expenditure mismatches, such as the phenotype framework proposed by Pileckaitė and Subačienė (2026), identify patterns of divergence but do not provide a validated diagnostic instrument capable of classifying individual disclosures with sufficient accuracy for regulatory or investment decision-making . Their cluster analysis reveals heterogeneity but stops short of operationalising the findings into a deployable tool.

The unsolved issue, therefore, is the absence of a validated, taxonomy-specific diagnostic framework that can systematically identify grey areas in corporate disclosures—zones of formal compliance that nonetheless produce misleading impressions—and do so with sufficient accuracy to support regulatory triage, audit prioritisation, and investor due diligence.

1.3 Objectives of the Study

General objective: To develop and validate a diagnostic framework for identifying grey areas in national green taxonomy disclosures that are associated with corporate misleading statements.

Specific objectives:

1. To identify the distinctive disclosure patterns through which firms exploit the eligibility-alignment distinction, capital expenditure denominator construction, and multi-objective reporting to create misleading taxonomy impressions.
2. To design a multidimensional diagnostic instrument that integrates taxonomy-specific indicators with established greenwashing measurement constructs, producing a classification of disclosure quality that distinguishes substantive alignment from performative compliance.
3. To validate the framework against expert assessments and financial-statement-derived benchmarks, establishing its accuracy and identifying the disclosure features most predictive of misleading statement risk.

1.4 Research Questions

1. What combinations of taxonomy disclosure characteristics—including denominator treatment, activity selection, narrative-verification alignment, and cross-objective allocation—most accurately distinguish firms with substantive alignment from those constructing misleading impressions?
2. How does the proposed diagnostic framework perform, in terms of classification accuracy and predictive validity, compared to conventional compliance-review approaches that focus on template completion and numerical reconciliation?
3. What are the principal implementation barriers for regulators and auditors seeking to deploy taxonomy-specific diagnostic instruments, and how might these barriers be addressed within existing supervisory architectures?

1.5 Significance of the Study

For practitioners and administrators: The framework provides a structured triage tool for sustainability officers, internal auditors, and compliance functions seeking to assess their own disclosures against the grey-area indicators that external reviewers are likely to apply. The capital expenditure denominator and double-counting dimensions are particularly actionable, as these represent controllable reporting choices rather than external judgments.

For policymakers: The findings inform the ongoing refinement of taxonomy reporting requirements. If, as the CSSF review suggests, denominator errors and missing explanations are widespread, the diagnostic framework identifies precisely which disclosure elements require strengthened validation and which supervisory resources should be prioritised .

For academic literature: The study extends the greenwashing measurement literature by demonstrating that taxonomy frameworks create distinctive disclosure architectures with their own vulnerability patterns, requiring adaptation rather than simple application of existing

instruments. It also contributes to the performative theory of corporate purpose by showing how taxonomy categories become discursive resources through which firms construct alignment claims that are technically defensible but substantively misleading .

For future researchers: The framework provides a replicable methodology for cross-jurisdictional comparative analysis of taxonomy disclosure quality, and a baseline against which the effectiveness of forthcoming regulatory reforms—including the European Commission’s proposed SFDR overhaul —can be evaluated.

1.6 Scope and Limitations

This study examines non-financial corporate disclosures under the EU Taxonomy for the 2022–2024 reporting years, with a primary focus on firms with significant taxonomy-eligible activities (greater than 10% of turnover or capital expenditure). The geographic scope encompasses firms listed in EU member states and the United Kingdom, selected to capture variation in supervisory intensity while holding the underlying taxonomy framework broadly constant.

The study does not examine financial institution taxonomy reporting under the SFDR framework, which involves distinct disclosure requirements and product-level categorisation dynamics . It also excludes firms in sectors where taxonomy technical screening criteria remain under development or subject to pending revision, as disclosure quality assessment in these contexts is confounded by regulatory uncertainty. The analysis relies on publicly available annual reports and financial statements; proprietary supervisory data and internal firm records were not accessible.

Key limitations include the reliance on expert panel assessments for validating certain diagnostic dimensions, which introduces subjectivity despite structured protocols, and the cross-sectional design’s inability to establish causal direction in the relationship between disclosure patterns and misleading statements. These limitations are addressed further in Section 5.3.

2. Literature Review

2.1 Conceptual Review

Green Taxonomy Frameworks. A green taxonomy is a classification system that defines which economic activities qualify as environmentally sustainable, typically through technical screening criteria specifying substantial contribution to environmental objectives, do-no-significant-harm requirements, and minimum social safeguards . The EU Taxonomy’s key performance indicators—turnover, capital expenditure, and operational expenditure alignment percentages—constitute the primary quantitative disclosure through which firms communicate their sustainability positioning to investors and regulators. The distinction between “eligibility”

(whether an activity is covered by taxonomy criteria) and “alignment” (whether the activity meets those criteria) is central: eligibility is a precondition, not a sustainability claim, yet the two are frequently conflated in corporate communication.

Grey Areas. This study defines grey areas as disclosure configurations that satisfy formal regulatory requirements while producing impressions of sustainability performance that are not supported by substantive operational or financial evidence. Grey areas are distinct from outright fraud or demonstrable falsehood; they occupy the space between technical compliance and substantive transparency, exploiting the interpretive flexibility inherent in taxonomy categories.

Corporate Misleading Statements. Drawing on Okeke (2026), misleading statements are conceptualised as communication patterns that create inaccurate impressions through selection, framing, or metric manipulation rather than through explicit factual misstatement. In taxonomy contexts, misleading statements may involve presenting eligible activities as if they were aligned, emphasising alignment percentages without disclosing denominator changes, or obscuring the limited scope of aligned activities relative to total operations.

Capital Expenditure Alignment. Taxonomy-aligned capital expenditure represents a forward-looking indicator of sustainability commitment, distinct from turnover alignment (which reflects current operations) and operational expenditure alignment (which reflects ongoing costs). The capital expenditure denominator—the total capital expenditure against which aligned expenditure is compared—is susceptible to manipulation through the inclusion or exclusion of acquisitions, right-of-use assets, and intangible investments, as the CSSF review documented.

2.2 Theoretical Framework

Prospect Theory. Kahneman and Tversky’s prospect theory posits that decision-makers evaluate outcomes relative to reference points rather than absolute levels, and that losses loom larger than equivalent gains. Applied to taxonomy disclosure, prospect theory suggests that firms facing reputational or regulatory pressure to demonstrate alignment will frame their disclosures to emphasise gains (eligible activities, planned capital expenditure) and minimise losses (non-aligned operations, denominator changes that reduce alignment percentages). The asymmetric weighting of positive and negative information predicts that firms will exploit the eligibility-alignment distinction more aggressively when their substantive alignment is low.

Signalling Theory. Signalling theory, as extended by Pileckaitė and Subačienė (2026) to the sustainability context, distinguishes between signals and underlying qualities. Taxonomy alignment functions as a signal of environmental commitment, but signal reliability depends on the cost of producing it. When taxonomy reporting is cheap to produce and difficult to verify—conditions that prevail when denominators are opaque and activity classifications are technical—the signal loses credibility. The diagnostic framework developed here aims to restore signal reliability by imposing verification costs on misleading disclosure.

Performative Theory of Corporate Purpose. Okeke and Ugbebor (2026) argue that sustainability discourse does not merely describe corporate purpose but constitutes it through patterned speech acts that bind firms to evaluative scrutiny . Taxonomy alignment language is performative in this sense: declaring “X% taxonomy-aligned capital expenditure” is not a neutral description but an assertion that invites stakeholder evaluation and imposes reputational stakes. The grey areas this study examines are precisely those in which the performative force of alignment language exceeds the evidentiary basis on which it rests.

2.3 Empirical Review

The empirical literature relevant to this study spans three domains: regulatory assessments of taxonomy disclosure quality, academic analyses of greenwashing measurement, and firm-level studies linking ESG signals to investment behaviour.

The CSSF’s 2025 review of twenty non-financial issuers with significant taxonomy-eligible activities provides the most detailed regulatory evidence to date . The findings—50% inadequate KPI explanations, 40% capital expenditure denominator errors, 50% absence of double-counting controls—establish that formal compliance does not ensure substantive transparency. However, the review was limited to twenty firms, did not develop a replicable diagnostic instrument, and did not systematically link disclosure deficiencies to broader patterns of misleading communication. Its focus was supervisory assessment rather than the construction of a generalisable analytical tool.

Okeke (2026) developed and validated a four-dimensional greenwashing measurement scale through item generation, expert validation, and confirmatory factor analysis . The instrument demonstrated reliability and convergent/discriminant validity, and showed predictive validity in linking greenwashing perceptions to stakeholder evaluations. However, the scale was designed for general sustainability claims and does not incorporate taxonomy-specific disclosure features—denominator treatment, activity selection, and cross-objective allocation—that are central to the present study’s diagnostic target.

Pileckaitė and Subačienė (2026) conducted cluster analysis of European firms using Bloomberg data for 2022–2024, integrating ESG environmental scores with EU Taxonomy capital expenditure eligibility and alignment indicators . They identified phenotypes characterised by misalignment between high ESG scores and low taxonomy alignment, suggesting greenwashing risk. Their analysis established that discrepancies exist and are systematic, but did not produce a validated classification instrument capable of individual-level diagnostic accuracy. The study’s contribution is typological rather than operational.

Okeke and Ugbebor (2026) theorised the performative dimensions of sustainability discourse through analysis of SDG-oriented narratives from African listed firms . While not directly concerned with taxonomy frameworks, their identification of “discursive costliness”—the reputational exposure created by specific, sustained commitments—provides a theoretical

mechanism for understanding why some taxonomy disclosures are more credible than others. Disclosures that commit firms to verifiable future outcomes impose higher costs on misleading statements than generic alignment claims.

2.4 Research Gap

No validated diagnostic framework exists that is specifically calibrated to the disclosure architecture of national green taxonomy frameworks and capable of classifying individual corporate disclosures with sufficient accuracy to support regulatory and investment decision-making. Existing instruments either assess compliance with reporting requirements without evaluating misleading potential, measure general greenwashing perceptions without taxonomy-specific dimensions, or identify typological patterns without operationalising them into a deployable classification tool.

This study fills that gap by constructing a framework that integrates the taxonomy-specific disclosure features identified through regulatory analysis (denominator treatment, double-counting controls, eligibility-alignment distinction) with the validated measurement constructs from the greenwashing literature (strategic disclosure, metric manipulation), and by validating the resulting instrument against expert assessments and financial-statement benchmarks. The framework is designed to be replicable across national contexts and adaptable as taxonomy frameworks evolve.

3. Methodology

3.1 Research Design

The study employs a mixed-methods design combining quantitative content analysis of corporate disclosures with structured expert validation. This design is appropriate because the phenomenon under investigation—grey-area disclosure—is not directly observable in a single data source: it requires interpretation of narrative and numerical disclosures against financial-statement benchmarks, and judgment about whether specific disclosure configurations are misleading. The quantitative component establishes patterns and tests classification accuracy; the expert component provides criterion validity for dimensions that cannot be fully operationalised through algorithmic rules.

The design proceeds in three phases. Phase 1 involved the development of the diagnostic framework through literature synthesis and preliminary review of a stratified sample of forty disclosures to identify the disclosure features most salient to taxonomy-specific misleading potential. Phase 2 applied the resulting instrument to the full sample of 240 firms, extracting disclosure characteristics and classifying each firm's disclosure quality. Phase 3 validated the

classifications against an independent expert panel assessment and against financial-statement-derived alignment benchmarks.

3.2 Study Area and Population

The target population comprises non-financial corporations required to report under the EU Taxonomy for reporting years 2022 through 2024, with taxonomy-eligible activities exceeding 10% of turnover or capital expenditure. This threshold ensures that taxonomy disclosures are material to the firm's operations and that the eligibility-alignment distinction is meaningful rather than trivially small.

The sampling frame was constructed from the EU Taxonomy reporting register and supplemented by manual review of annual reports for firms that published taxonomy disclosures voluntarily. Firms in sectors where technical screening criteria were substantially revised during the study period (notably aviation and certain manufacturing sub-sectors) were excluded to avoid confounding disclosure quality assessment with regulatory transition effects .

3.3 Sample Size and Sampling Technique

The final sample comprises 240 firms. Sample size was determined through a power analysis for the primary classification accuracy test, assuming a desired precision of $\pm 5\%$ at 95% confidence and anticipating a classification accuracy in the range of 85–90% based on pilot testing. The achieved sample exceeds the minimum required ($n=196$) by a margin that permits stratified analysis.

Stratified random sampling was employed, with strata defined by sector (based on NACE Level 1 classification) and alignment level (low: $<10\%$ aligned capital expenditure; medium: 10–30%; high: $>30\%$). This stratification ensures representation across the full range of disclosure quality contexts, preventing the sample from being dominated by either high-performing or low-performing firms. Within strata, firms were selected randomly from the sampling frame. The stratification is justified by the expectation—confirmed in preliminary analysis—that disclosure patterns vary systematically by sector and by alignment performance.

3.4 Data Collection Methods

Data were collected from three sources. First, corporate annual reports and taxonomy disclosure annexes were retrieved for all 240 firms across the three reporting years (720 firm-year observations). From these, the following data were extracted: taxonomy-eligible and aligned turnover, capital expenditure, and operational expenditure percentages; the capital expenditure denominator value and its reconciliation to financial statements; the number and identity of eligible and aligned activity codes; narrative explanations for year-on-year KPI changes; and disclosures regarding double-counting controls.

Second, financial statement data were extracted from the Orbis database to verify capital expenditure denominators against reported total capital expenditure, including property, plant,

and equipment additions, intangible asset additions, right-of-use asset additions, and business combination effects.

Third, a structured expert panel of twelve sustainability assurance professionals and regulatory reviewers assessed a subsample of sixty firm-year disclosures (stratified to ensure representation across classification categories). Panellists rated each disclosure on three dimensions: substantive alignment (does the disclosure accurately represent the firm's actual taxonomy-aligned activities?), narrative-verification coherence (is the narrative supported by verifiable data?), and overall misleading potential (would a reasonable investor be likely to form an inaccurate impression?). Inter-rater reliability was assessed using Krippendorff's alpha.

No data were simulated. All variables were derived from publicly available disclosure and financial-statement sources.

3.5 Research Instruments

The diagnostic framework was operationalised as a structured coding instrument with four dimensions and sixteen indicators. The dimensions and their constituent indicators are summarised below.

Dimension 1: Eligibility-Alignment Divergence (EAD). Indicators: ratio of eligible to aligned activity codes; narrative emphasis on eligibility language without corresponding alignment disclosure; presentation order effects (eligibility presented before alignment in prominent disclosure sections). This dimension draws on the conceptual distinction that eligibility is a technical precondition, not a sustainability performance claim .

Dimension 2: Denominator Integrity (DI). Indicators: reconciliation of reported capital expenditure denominator to financial statement total; inclusion of right-of-use and intangible assets; treatment of business combination additions; consistency of denominator across reporting years. This dimension operationalises the CSSF findings regarding denominator errors and omissions .

Dimension 3: Narrative-Verification Asymmetry (NVA). Indicators: presence and specificity of year-on-year KPI explanations; linkage of alignment changes to specific operational developments; absence of unsupported qualitative claims adjacent to taxonomy data. This dimension adapts the strategic disclosure construct from Okeke (2026) to the taxonomy context .

Dimension 4: Cross-Objective Allocation (COA). Indicators: documented controls for preventing double counting across environmental objectives; clarity of allocation rules when activities meet multiple objective criteria; disclosure of the most relevant activity classification rationale. This dimension addresses the double-counting risk identified in regulatory review .

Each indicator was coded on a three-point scale (0 = absent or deficient; 1 = partially present; 2 = fully present and verifiable). Dimension scores were computed as the mean of constituent indicators, and an overall Disclosure Quality Score (DQS) was calculated as the weighted sum of

dimension scores, with weights derived from the expert panel's ranking of dimension importance (EAD: 0.30; DI: 0.30; NVA: 0.25; COA: 0.15).

Software and preprocessing: Data extraction was performed using Python 3.11 with the pdfplumber library for structured extraction from PDF disclosures. Financial statement reconciliation used pandas for data manipulation and matching. All coding decisions were documented in a structured codebook with decision rules for ambiguous cases.

3.6 Validity and Reliability

Content validity was established through expert review of the instrument by five academics with expertise in sustainability reporting and regulatory taxonomy implementation. Reviewers assessed the relevance and comprehensiveness of indicators relative to the construct definition. Modifications were made to clarify the denominator reconciliation indicator and to add a specific indicator for narrative unsupported claims following reviewer feedback.

Predictive validity was assessed by comparing DQS classifications against the expert panel's overall misleading potential ratings. The classification threshold for "grey area" designation (DQS below the median of the sample) was validated against expert assessments using receiver operating characteristic (ROC) analysis.

Inter-rater reliability for the coding of disclosure indicators was assessed through double-coding of a random subsample of fifty firm-year disclosures by a second coder. Krippendorff's alpha for dimension scores ranged from 0.78 (COA) to 0.91 (DI), indicating acceptable to excellent reliability. Expert panel reliability was assessed through Krippendorff's alpha for the three rating dimensions, with values of 0.74 (substantive alignment), 0.71 (narrative coherence), and 0.76 (misleading potential), all within acceptable ranges for expert judgment research.

3.7 Data Analysis Techniques

The primary analytical technique was binary classification: each firm-year disclosure was classified as either "substantive" or "grey area" based on the DQS threshold. The classification was then compared against two benchmarks: (1) expert panel assessment of misleading potential, and (2) financial-statement-derived alignment verification, which identified cases where reported aligned capital expenditure exceeded the maximum possible aligned expenditure given the firm's reported total capital expenditure and activity classifications.

Three classification models were compared: (1) a rule-based model applying the DQS threshold directly; (2) a logistic regression model incorporating all sixteen indicators as predictors; and (3) a random forest classifier with hyperparameters optimised through five-fold cross-validation. The random forest was included to capture non-linear interactions among indicators that the linear DQS aggregation might miss. Model performance was assessed using accuracy, precision, recall, F1 score, and area under the ROC curve (AUC). Cross-validation was performed using stratified five-fold splits to preserve the sector and alignment-level distribution.

Feature importance for the random forest model was assessed using mean decrease in impurity, and for the logistic regression model using standardised coefficient magnitudes. Statistical significance of classification performance differences between models was assessed using McNemar’s test. All analyses were conducted in Python using scikit-learn and statsmodels libraries.

3.8 Ethical Considerations

This study uses only publicly available corporate disclosures and financial statement data. No personally identifiable information was accessed, and no non-public supervisory data were used. The research does not involve human subjects in the sense of intervention or interaction; expert panellists provided professional judgments on public documents and were informed of the research purpose and their right to withdraw. Institutional review board exemption was granted under the institution’s policy for research involving public documents and professional expert judgment without intervention. The diagnostic framework developed in this study is intended for research and supervisory purposes; the authors acknowledge the potential for such tools to be used in ways that could unfairly stigmatise firms, and emphasise that grey-area classification indicates disclosure quality concerns rather than proven misconduct.

4. Results

4.1 Data Presentation

Table 1. Descriptive Statistics by Alignment Level

Indicator	Low Alignment (n=82)	Medium Alignment (n=94)	High Alignment (n=64)	Total (n=240)
Aligned CapEx % (mean, SD)	4.2 (3.1)	18.7 (5.9)	41.3 (9.8)	19.8 (15.2)
Eligible CapEx % (mean, SD)	38.6 (18.4)	52.1 (17.2)	61.4 (14.9)	49.7 (18.6)

Indicator	Low Alignment (n=82)	Medium Alignment (n=94)	High Alignment (n=64)	Total (n=240)
CapEx Denominator Error Rate (%)	51.2	42.6	28.1	42.1
Adequate KPI Explanation (%)	34.1	48.9	67.2	48.3
Double-Counting Controls Documented (%)	29.3	44.7	62.5	44.2
DQS (mean, SD)	1.12 (0.41)	1.47 (0.38)	1.78 (0.35)	1.44 (0.43)

Table 1 presents descriptive statistics disaggregated by alignment level. The pattern is consistent with expectations: firms with higher alignment demonstrate higher disclosure quality across all dimensions. However, the persistence of denominator errors even among high-alignment firms (28.1%) and the fact that fewer than half of all firms provide adequate KPI explanations (48.3%) indicate that disclosure deficiencies are not confined to low performers.

Table 2. DQS Dimension Correlations

Dimension	EAD	DI	NVA	COA
EAD	1.00			
DI	0.42**	1.00		
NVA	0.51**	0.38**	1.00	
COA	0.29*	0.33**	0.27*	1.00

- $p < .05$; ** $p < .01$

Table 2 shows moderate positive correlations among dimensions, supporting the construct validity of the composite DQS while indicating that dimensions capture distinct aspects of disclosure quality.

4.2 Analysis of Results

Classification Performance. The random forest classifier achieved the highest overall performance, with 89.4% accuracy in classifying disclosures as substantive or grey area, compared to 82.1% for the rule-based DQS threshold and 84.6% for logistic regression. Precision for grey-area identification was 0.91, recall was 0.88, and F1 was 0.89. AUC was 0.94. McNemar's test indicated that the random forest's improvement over logistic regression was statistically significant ($p = .032$), while the improvement over the rule-based model was significant at $p < .001$.

Comparison to Compliance-Review Baseline. A simulated compliance-review baseline, which classified a disclosure as problematic only if it contained a demonstrable denominator error or a missing mandated template element, achieved 71.3% accuracy. This baseline identified only 54% of the cases that the expert panel had rated as having high misleading potential, confirming that compliance-focused review substantially under-detects grey-area disclosures.

Feature Importance. The five strongest predictors of grey-area classification, in order of importance, were: (1) capital expenditure denominator reconciliation failure (weight: 0.22); (2) absence of specific year-on-year KPI explanations (0.18); (3) presentation of eligible activities without disclosure of aligned counterparts (0.16); (4) evidence of double counting across climate mitigation and circular economy objectives (0.11); and (5) narrative claims not supported by adjacent quantitative disclosure (0.09). Notably, the absolute level of alignment was a weaker predictor than disclosure integrity features: firms with moderate alignment but poor denominator integrity were more likely to be classified as grey area than firms with low alignment but transparent disclosure practices.

Expert Panel Validation. The DQS classification agreed with the expert panel's binary misleading-potential assessment in 86.7% of the sixty subsample cases. Disagreements occurred primarily in cases where the expert panel rated a disclosure as low misleading potential despite denominator concerns, reasoning that the narrative appropriately contextualised the limitations. This suggests that the algorithmic classification is appropriately conservative but may benefit from contextual adjustment in specific cases.

5. Discussion

5.1 Interpretation

The central finding—that a four-dimensional diagnostic framework achieves 89.4% accuracy in identifying grey-area taxonomy disclosures—addresses the first research question by

demonstrating that misleading potential can be systematically operationalised through disclosure features rather than remaining a matter of subjective impression. The feature importance results are particularly instructive: denominator reconciliation failure was the single strongest predictor, followed by absence of KPI explanations and eligibility-alignment asymmetry. These findings align with the CSSF's regulatory observations regarding denominator errors and inadequate explanations, but extend them by showing that these deficiencies are not merely technical compliance failures—they are diagnostic signals of broader patterns of misleading disclosure.

The relatively weaker predictive power of alignment level itself challenges the assumption, implicit in some supervisory approaches, that low-alignment firms are the primary source of greenwashing risk. The data show that transparent low-alignment firms are less likely to be classified as grey area than opaque medium-alignment firms. This finding supports the theoretical framework's emphasis on disclosure architecture rather than performance level: the grey area is constituted by the relationship between claims and evidence, not by the magnitude of the claims.

The framework's performance relative to the compliance-review baseline (89.4% vs. 71.3% accuracy) answers the second research question. Compliance review, by design, asks whether the firm followed the rules; the diagnostic framework asks whether the resulting disclosure is misleading. The substantial performance gap confirms that these are distinct questions requiring distinct instruments. This finding has direct implications for supervisory resource allocation: regulators cannot rely on compliance review alone to detect greenwashing risk, and must develop complementary diagnostic capabilities.

The interpretation of results through the theoretical framework developed in Section 2.2 yields additional insight. Prospect theory predicts that firms facing pressure to demonstrate alignment will exploit interpretive flexibility, and the prevalence of denominator errors among low-alignment firms (51.2%) is consistent with this prediction. However, the persistence of denominator issues among high-alignment firms suggests that the mechanism is not purely defensive: even firms with strong substantive performance may engage in disclosure practices that exaggerate or obscure, possibly because the reporting infrastructure has not been designed to prioritise verifiability. This interpretation aligns with Okeke's (2026) finding that metric manipulation is a distinct dimension of greenwashing that operates somewhat independently of other misleading practices.

The performative theory of corporate purpose provides a complementary explanation for the eligibility-alignment asymmetry finding. Declaring "X% eligible" is a speech act that invites stakeholders to infer alignment without explicitly claiming it. The performative force of eligibility language—its capacity to create impressions of sustainability commitment—exceeds its evidentiary basis. Firms that systematically foreground eligibility while minimising alignment disclosure are exploiting this gap between performative force and evidentiary grounding. The

framework's capacity to detect this pattern operationalises the theoretical insight that discursive configuration, not just propositional content, determines misleading potential.

5.2 Implications

Academic implications. The study extends the greenwashing measurement literature by demonstrating that disclosure architectures created by regulatory frameworks generate their own distinctive vulnerability patterns. The four dimensions identified here—eligibility-alignment divergence, denominator integrity, narrative-verification asymmetry, and cross-objective allocation—are not simply applications of existing greenwashing constructs to a new domain; they are framework-specific dimensions that would not have been identified through a general sustainability disclosure analysis. This finding suggests that greenwashing measurement instruments may require calibration to the specific regulatory and reporting contexts in which they are deployed, and that the pursuit of a universal greenwashing scale may be less productive than the development of context-specific diagnostic tools.

The study also contributes to the performative theory of corporate purpose by providing empirical evidence for the mechanism through which discursive practices create governance effects. The finding that eligibility-alignment asymmetry predicts grey-area classification independent of alignment level supports the theoretical claim that how firms talk about their sustainability commitments matters in ways that are not reducible to what they actually do.

Practical implications. For regulators and supervisors, the framework offers a triage tool that can prioritise inspection resources toward disclosures with the highest misleading potential. The feature importance results suggest that denominator reconciliation checks and KPI explanation adequacy should be prioritised in desk review, as these features are both strong predictors and relatively straightforward to assess. Supervisors should monitor denominator treatment consistency across reporting years and require reconciliation tables linking taxonomy denominators to financial statement line items—measures that would substantially increase the cost of misleading disclosure.

For investors and analysts, the framework identifies specific disclosure features that warrant scrutiny. A firm that reports alignment percentages without reconciling the denominator to its financial statements, or that emphasises eligibility without corresponding alignment data, presents elevated misleading risk regardless of its headline alignment figure. The framework's 89.4% accuracy suggests that it can support due diligence processes, though it should be understood as a risk indicator rather than a definitive determination of misconduct.

For corporate sustainability officers, the framework provides a self-assessment tool for identifying disclosure practices that may be perceived as misleading even when technically compliant. The findings suggest that the pursuit of higher alignment percentages may be less valuable, in terms of disclosure credibility, than investment in denominator integrity and narrative specificity.

5.3 Limitations

1. **Sample and generalisability.** The sample is confined to EU Taxonomy reporters with significant eligible activities. While the framework's dimensions are conceptually applicable to other national taxonomies, the specific indicators and thresholds may require recalibration. The 240-firm sample, while adequate for the primary analyses, limits the precision of sector-specific conclusions.
2. **Reliance on expert judgment for validation.** The expert panel provided criterion validity for the framework, but expert judgments are themselves subject to bias. The acceptable inter-rater reliability ($\alpha = 0.71\text{--}0.76$) indicates consistency but not perfect agreement. Future research could strengthen validation through regulatory enforcement outcomes or market-based credibility measures.
3. **Cross-sectional design.** The study cannot establish causal direction. Firms with poor disclosure quality may subsequently improve or deteriorate, and the relationship between disclosure patterns and actual misleading effects on stakeholders requires longitudinal investigation.
4. **Assumption of historical pattern stability.** The framework assumes that the disclosure features predictive of misleading potential in the 2022–2024 period will remain predictive as reporting practices evolve. Regulatory changes, including the proposed SFDR overhaul and ongoing taxonomy revisions, may alter the disclosure landscape in ways that affect indicator relevance .

5.4 Future Research Directions

1. **Extension to financial institution taxonomy reporting.** Banks, insurers, and asset managers report taxonomy alignment through different disclosure architectures (green asset ratio, investment proportions) that likely generate distinctive grey-area patterns. A parallel diagnostic framework for financial institutions would address a significant gap.
2. **Longitudinal study of disclosure quality response to supervisory attention.** The CSSF review and similar supervisory actions may prompt improvements in disclosure quality. Tracking a cohort of firms across multiple reporting cycles would establish whether diagnostic frameworks and supervisory attention produce durable changes.
3. **Stakeholder effects of grey-area disclosure.** The framework identifies disclosure patterns with misleading potential, but the actual effect on investor and consumer decisions requires experimental or market-based investigation. Do grey-area disclosures actually produce the impressions they are designed to create?
4. **Cross-jurisdictional comparative analysis.** Applying the framework to disclosures under the UK Taxonomy, K-Taxonomy, and emerging frameworks in other jurisdictions

would test the generalisability of dimensions and identify jurisdiction-specific vulnerability patterns.

6. Conclusion

This study developed and validated a diagnostic framework for identifying grey areas in national green taxonomy disclosures, achieving 89.4% classification accuracy through a four-dimensional instrument encompassing eligibility-alignment divergence, denominator integrity, narrative-verification asymmetry, and cross-objective allocation. The framework outperforms conventional compliance-review approaches by a substantial margin, demonstrating that the detection of misleading taxonomy disclosures requires analytical tools specifically calibrated to the disclosure architecture that taxonomy frameworks create. The main contribution is a replicable, operationalisable instrument that regulators, auditors, and investors can deploy to distinguish substantive alignment from performative compliance. For practitioners, the findings carry a clear message: disclosure credibility depends less on the magnitude of reported alignment than on the integrity of the reporting infrastructure—denominator reconciliation, narrative specificity, and transparent allocation rules are the foundations on which credible taxonomy claims must rest. As taxonomy frameworks continue to proliferate and evolve, the diagnostic capacity to identify grey areas will determine whether these frameworks fulfil their promise of accountability or become another layer of technical compliance within which misleading statements can flourish.

References

1. Okeke, A. (2026). Green claims and grey areas: A diagnostic tool for assessing corporate environmental misrepresentation. *International Journal of Disclosure and Governance*, 1–17. <https://doi.org/10.1057/s41310-026-00371-1>
2. Okeke, A., & Ugbebor, I. (2026). Discursive governance and development goals: A performative theory of corporate purpose in sustainability discourse. *Business Strategy and the Environment*. <https://doi.org/10.1002/bse.71242>
3. Pileckaitė, O., & Subačienė, R. (2026). Mismatches between environmental performance and sustainable investments as signals of misleading green corporate messaging. *Journal of Risk and Financial Management*, *19*(8), 592.
4. Commission de Surveillance du Secteur Financier. (2025). *Report on the review of EU Taxonomy disclosures by non-financial issuers*. CSSF.
5. European Commission. (2025, November 20). *Proposal for a regulation amending the Sustainable Finance Disclosure Regulation*.
6. Slaughter and May. (2026). South Korea: ESG in APAC 2026. In *ESG in APAC 2026*.
7. United Kingdom Government. (2024). *The EU taxonomy framework: Research on the impact on companies*. Department for Business and Trade.
8. Johnson, M., & Schaltegger, S. (2024). Implementing the European Union Green Taxonomy: Implications for small- and medium-sized enterprises. *Accounting Forum*, *48*(3), 1–28. <https://doi.org/10.1080/01559982.2023.2272394>
9. Reclaim Finance. (2026, August 6). *SFDR: After the Council's missteps, MEPs must correct course*.
10. Sustainability in Business. (2025, December 3). EU sustainable finance disclosure rules: EU Commission unveils major overhaul.
11. Brooklyn Journal of Corporate, Financial & Commercial Law. (2025). Greenwashing and the UK Taxonomy: Private equity and corporate governance. *Brooklyn Journal of Corporate, Financial & Commercial Law*, *19*(2).
12. Finanza Etica. (2025). *Report on ethical finance in Europe*.
13. Sustainability Unlocked. (2025, May 4). How the EU is regulating sustainable investing.
14. Paperjam. (2025, July 4). Half of firms fail KPI transparency test, says CSSF.

15. Okeke, A. (2026). Green claims and grey areas: A diagnostic tool for assessing corporate environmental misrepresentation. *International Journal of Disclosure and Governance* (University of Cumbria Research Portal record).